

Liquid-cooled CRAY T3E systems

Scalable power that works

The CRAY T3E series of scalable parallel systems is the most powerful and flexible in the world. Across a wide range of air- and liquid-cooled configurations from 16 to 2048 processors with corresponding peak performance levels from 9.6 GFLOPS to 1.2 TFLOPS and prices from under \$1 million, CRAY T3E systems deliver unmatched scalable performance and price/performance.

The power and flexibility of this new generation come from global system scalability—

scalable performance based on the world's fastest microprocessor and the industry's best interprocessor communications, the world's first scalable operating system, and scalable I/O. In any configuration, every function and feature contributes to balanced scalable performance. Every price point offers tremendous room for growth with a corresponding payback in performance, throughput, and productivity.

CRAY T3E scalable parallel systems use the DECchip 21164 (DEC Alpha EV5) from Digital Equipment Corporation, capable of 600 MFLOPS peak performance. Each CRAY T3E processing element (PE) has its own local DRAM memory with a capacity of from 64 Mbytes to 2 Gbytes. A shared, high-performance, globally addressable memory subsystem makes these memories accessible to every PE in a CRAY T3E system.

To support the scalability of the CRAY T3E system, we created UNICOS/mk, a scalable version of our UNICOS operating system. UNICOS/mk is divided into "servers," which are distributed among the

Performance


- 64 to 2048 processing elements
- 38.4 GFLOPS - 1.2 TFLOPS peak performance
- Up to 4 Tbytes of globally addressable memory, with 2 Gbytes per PE memory option
- High-performance interconnect/3-D torus
- High-bandwidth, scalable I/O, up to 250,000 Mbytes/s (250 Gbytes/s)

User productivity—complete, easy-to-use programming environment

- UNICOS/mk serverized, microkernel-based operating system
- Scalable, distributed functionality
- Standard, open UNIX operating environment
- Systemwide suite of resources managed as single entity
- CRAY CF90 programming environment
- CRAFT and HPF programming models
- CRAY C++ programming environment, including C
- Highly optimized Cray Message Passing Toolkit
 - PVM
 - MPI
- Optimized scientific libraries
- Cray TotalView debugger
- Advanced performance analysis tools
- Visual program browser
- Automatic data conversion (CRAY, IBM, CDC)
- Message passing to support development of scalable heterogeneous applications

Production capabilities

- Dynamic scheduling of batch and interactive processes
- Large job processing
- Efficient parallel processing
- Full connectivity to UNIX (TCP/IP) and proprietary networks
- Security features (MLS, Kerberos, RIPS0, CIPSO)
- Extensive accounting and resource management facilities
- Network Queuing Environment
- Automated hierarchical storage manager
 - Data Migration Facility (DMF)
- Tape volume management system
- Online tape support
- STK support
- Disk striping



processors of a CRAY T3E system. Local servers process operating system requests specific to each user PE, while global servers provide systemwide operating system capabilities. The result is the world's first scalable operating system. This distribution of functions provides a global view of the computing environment—a single-system image—that allows administrators to manage a systemwide suite of resources as a single entity.

Breakthrough I/O technology sets new standards for I/O scalability, performance, capacity, and configurability. The CRAY T3E system performs I/O through multiple ports onto one or more scalable GigaRing channels. Each dual-ring I/O channel, with data in the two rings traveling in opposite directions, delivers high I/O data bandwidth and enhances reliability. Liquid-cooled CRAY T3E systems may be configured with an I/O channel for every 16 PEs. All I/O channels are accessible and controllable from all PEs. On the CRAY T3E system, each GigaRing channel has a maximum data payload bandwidth of 1 Gbyte/s and provides high-speed access to peripherals, networks, and other Cray Research systems.

Liquid-cooled CRAY T3E scalable parallel systems are available in configurations from 64 to 2048 processors with peak performance levels from 38.4 GFLOPS to 1.2 TFLOPS.

As an option, CRAY T3E systems may be coupled closely to the parallel vector architecture of industry-leading CRAY T90 or CRAY J90 systems in a scalable heterogeneous system. This capability provides maximum flexibility to meet the needs of individual user workloads. Customers can choose from a wide variety of system solutions with different mixes of parallel vector and scalable parallel computational capabilities—an offering unique to Cray Research.

Liquid-cooled CRAY T3E system specifications

CRAY T3E system	
Peak performance	38.4 GFLOPS - 1.2 TFLOPS
Architecture	Multiple instruction multiple data (MIMD) with hardware support for single instruction multiple data (SIMD) processing
Processing Elements (PEs)	
Microprocessor	DECchip 21164 64-bit super-scalar RISC Four-way instruction issue with two floating-point operations per clock 32-bit and 64-bit IEEE floating-point arithmetic
Local memory per PE	64, 128, 256, or 512 Mbytes or 1 Gbyte or 2 Gbytes
Data error protection	SECCED
Clock speed	300 MHz
Peak performance per PE	600 MFLOPS
Peak memory bandwidth	1.2 Gbytes/s
Packaging	8 PEs per module
PEs per system	64 to 2048 user PEs in increments of 8 PEs. Additional support PEs are provided for operating system use and to provide resiliency
Memory	
Technology	16-Mbit or 64-Mbit DRAM
Architecture	Shared, physically distributed, globally addressable
Total system memory	4 Gbytes - 4 Tbytes
Interconnect network	
Topology	3-D torus
Data payload bisection bandwidth (512 PEs)	122 Gbytes/s
I/O	
Peak I/O bandwidth	1 I/O channel available for every 16 PEs 1 Gbyte/s per channel
Physical characteristics	
Cabinet footprint area	LC256: 35.4 ft ² (3.2 m ²) LC512: 57.4 ft ² (5.2 m ²) LC1024: 114.8 ft ² (10.4 m ²) LC2048: 229.6 ft ² (20.8 m ²)

Liquid-cooled CRAY T3E system configurations

Model ¹	Cabinets	User PEs	Base Memory ² (Gbytes)	Peak Performance (GFLOPS)	Maximum I/O Channels	I/O Bandwidth ³
LC64	1	64	8	38.4	4	4
LC128	1	128	16	76.8	8	8
LC256	1	256	32	150	16	16
LC384	2	384	48	230	24	24
LC512	2	512	64	300	32	32
LC1024	4	1024	128	600	64	64
LC2048	8	2048	256	1.2 TFLOPS	128	128

1. This is not a complete list of all liquid-cooled CRAY T3E systems. Any configuration above 64 PEs is possible, in increments of 8 PEs.
2. Base model configuration includes 128 Mbytes of memory per PE. Optional memory sizes of 64, 256, or 512 Mbytes or 1 Gbyte or 2 Gbytes per PE available.
3. The maximum I/O bandwidths for liquid-cooled systems may be doubled by special order.

CRAY and UNICOS are federally registered trademarks, and CRAY J90, CRAY T90, CRAY T3E, GigaRing, and Network Queuing Environment are trademarks of Cray Research, Inc. DECchip is a trademark of Digital Equipment Corporation. The Cray TotalView debugger is based on the TotalView debugger from Bolt, Beranek and Newman Inc. TotalView is a trademark of Bolt, Beranek and Newman Inc. Research for the development of the CRAY T3E system and the GigaRing channel was partially supported by the Advanced Research Projects Agency under Agreement #MDA972-95-3-0032, dated 12 January 1995. ESCON is a trademark of International Business Machines Corporation. Kerberos is a trademark of the Massachusetts Institute of Technology. UNIX is a registered trademark in the United States and other countries, licensed exclusively through X/Open Company Limited. X/Open is a trademark of X/Open Company Limited.